Washington Systems Center - Storage

A9000 / A9000R and SVC Best practices

Steve Solewin – IBM Corporate Solutions Architect

Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2018...

Washington Systems Center – Storage experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located in the Accelerate with IBM Storage Blog:
https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/playlist?list=PLSdmGMn4Aud-gKUBCR8K0kscCiF6E6ZYD&disable_polymer=true

2018 Webinars:

January 9 – DS8880 Easy Tier
January 17 – Start 2018 Fast! What's New for Spectrum Scale V5 and ESS
February 8 - VersaStack - Solutions For Fast Deployments
February 16 - TS7700 R4.1 Phase 2 GUI with Live Demo
February 22 - DS8880 Transparent Cloud Tiering Live Demo
March 7 - Spectrum Storage Management, Control, Insights, Foundation; what's the difference?
March 15 - IBM FlashSystem A9000/R and SVC Configuration Best Practices
Register Here: https://ibm2.webex.com/ibm2/onstage/g.php?MTID=e87d423c5cccbdefbbb61850d54f70f4b
March 27 - IBM FlashSystem A9000/R Technical Update
Register Here: https://ibm2.webex.com/ibm2/onstage/g.php?MTID=eb3bcbb054f1aabb189ea53cf2e8d7926

IBM Block Storage Virtualization Strategy - *Spectrum Virtualize*

Manage the storage pool from a central point

Virtual Disk
Virtual Disk
Virtual Disk
Virtual Disk

Make changes to the storage without disrupting host applications

SAN

Automated use of FlashSystem/SSD through tiering with Easy Tier

Spectrum Virtualize

Enable Real Time Compression, Thin Provisioning, Encryption and Transparent Cloud Tiering for all storage

Internal disk
Flash
IBM
HP

Storage Pool

Oracle
HDS
EMC
NetApp

Apply common Advanced Copy Services and HyperSwap across the storage pool

Combine the capacity from multiple arrays into a single pool of storage

SVC Logical Configuration View

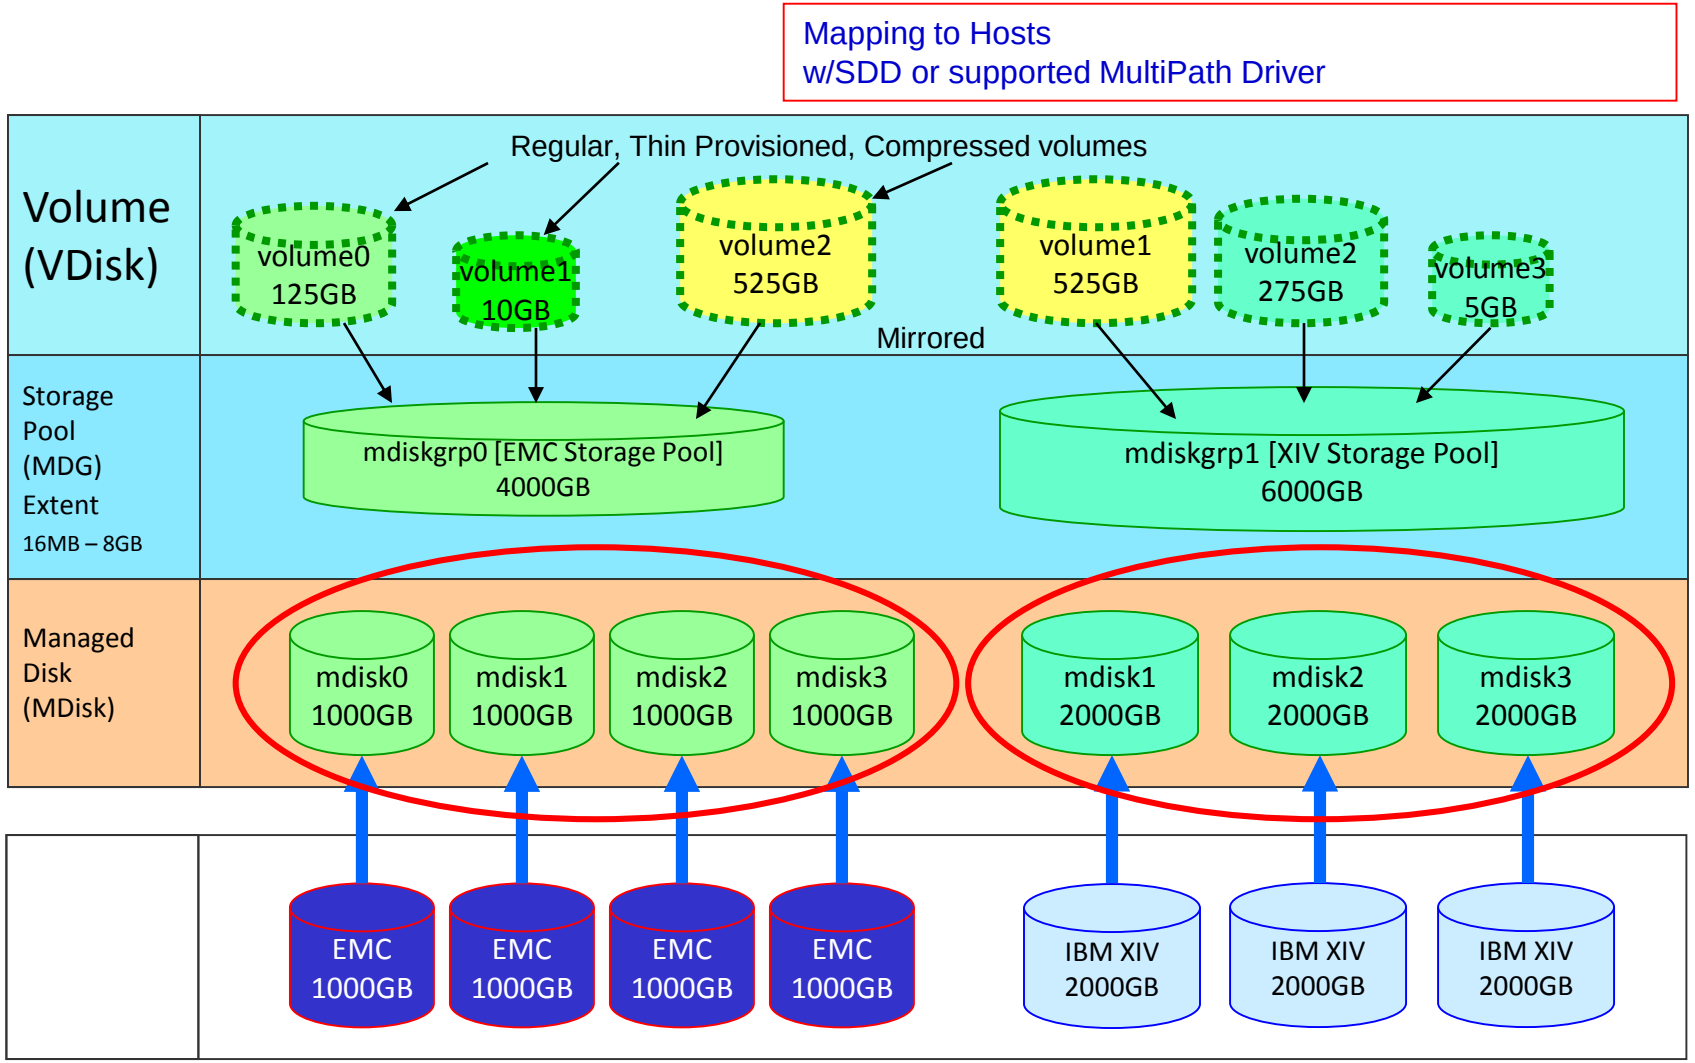

SVC Features and Functions

- **Current SVC hardware SV1 with 8.1.x code**
- **SVC features and functions**
  - Real-time compression
  - Transparent Cloud Tiering (TCT)
  - HyperSwap
  - FCP (16Gb), FCoE and iSCSI host support
  - High Availability and multi-site mirroring configurations
    - Synchronous and asynchronous mirroring
    - Volume (VDisk) mirroring
  - Native IP replication
  - Snapshots
  - Enhanced Stretched Cluster
  - Consistency group support
  - Thin provisioning
  - Online data migration
  - Easy Tier support
    - 3 Tier support
    - Auto rebalance
    - Warm Promote

Spectrum Virtualize 7.8 and 7.8.1

- **7.8 Updates**
  - Transparent Cloud Tiering
  - Tier 1 SSD's with Easy Tier
  - 92F Dense Drawer
  - Model Conversions
  - Support for SKLM
  - Software support for SCU licensing

- **7.8.1 Updates**
  - QoS Improvements
  - Increased FC background copy rates
  - Host Cluster GUI support
  - Consistency protection for Metro/Global Mirror
  - Resize VDisks in MM/GM relationship
  - Proactive Host failover

Spectrum Virtualize 8.1

- **Performance**
  - Dual socket multi threading
  - Significant improvement in (SV1 engine)
    - 4K Read Hit/Miss
    - 8K Read Hit/Miss
    - 4K 70/30 RW Miss
    - 8K 70/30 RW Miss
- **Support for up to 256GB memory**
- **Hot Spare node**
- **Updated GUI**
- **Remote Support Assistance**
- **Health Checker**
- **Up to 4 SKLM key servers**

https://www.ibm.com/support/knowledgecenter/STPVGU

Spectrum Virtualize 8.1.2

- Reduce costs with new data reduction pools
  - SCSI UNMAP support for deallocation and reclamation of unused capacity
  - Consistent and predictable performance for compressed volumes with new log-structured pool implementation enables more widespread use
  - Support for more compressed volumes: all volumes in a system (**up to 10,000** volumes)
- **Up to 2x** workload bandwidth for data intensive applications such as analytics
  - 25GbE HBA with RoCE protocol
  - 25GbE HBA with iWARP protocol
  - Supported with SVC Model SV1, StorwizeV7000F, V7000 Gen 2+, StorwizeV5030F, V5000 Gen 2, FlashSystem V9000 Model AC3

https://www.ibm.com/support/knowledgecenter/STPVGU

SVC / A9000 / A9000R code prereqs

- **Minimum 7.4 version 7.4.0.10**
- **Minimum 7.5 version 7.5.0.8**
- **Minimum 7.6 version 7.6.1.4**
- **Version 7.7 and higher**
- **Version 8.1 and higher**
- **Ensure the SVC is upgraded BEFORE attaching A9000/A9000R**

Pool Configuration Considerations

- **Typically recommend creating only one SVC storage pool (managed disk group) per A9000/R system**
  - Simplicity of management
  - Use multiple pools for special considerations
    - Workload Isolation – still be logical vs physically separated with A9000
    - If share nothing workload scenarios are necessary (i.e. SVC and direct A9000 connectivity)
  - SVC write cache is partitioned by storage pool

Number of SVC Storage Pools	Maximum occupancy allowance as % of cache
1	100
2	75
3	40
4	30
5 or more	25

Max Limitations and Best Practice recommendations

- **SVC Limits and Restrictions, Supported Hardware, and Product Documentation URL**
  - http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003658

Generic Host Properties			
Host objects (IDs) per system	2145-CF8 2145-CG8 2145-DH8 2145-SV1	2048	A host object may contain both Fibre Channel ports and iSCSI names
Host objects (IDs) per I/O group	2145-CF8 2145-CG8 2145-DH8 2145-SV1	512	Refer to the additional Fibre Channel and iSCSI host limits below
Volume mappings per host object	2145-CF8 2145-CG8 2145-DH8 2145-SV1	2048*	* Although SVC allows the mapping of up to 2048 volumes per host object, not all hosts are capable of accessing/managing this number of volumes. The practical mapping limit is restricted by the host OS, not SVC. Note: this limit does not apply to hosts of type adminlun (used to support VMware vvols).
Total Fibre Channel ports and iSCSI names per system		8192	
Total Fibre Channel ports and iSCSI names per I/O group		2048	
Total Fibre Channel ports and iSCSI names per host object		32	
iSCSI names per host object (ID)		8	

SVC/XIV-A9000/R Considerations

- **SVC Extent size guideline**
  - SVC supports storage pool extent sizes of 16, 32, 64, 128, 256, 512, 1024, 2048, and 8192 MB
    - Property of the SVC Storage Pool definition
    - Affects capacity SVC can manage vs. performance
    - Recommend 1 GB extent size for XIV/A9000
      - Higher data reduction rates may require 2GB extent

- **Identifying an A9000/R LUN in the SVC**
  - When an A9000/R volume is mapped to SVC, the A9000/R LUN ID is displayed in Decimal
  - When SVC discovers it, SVC brings it in as the corresponding Hex value which is placed into the field "Controller LUN Number"
    - Example, if the A9000/R GUI shows LUN 18, the SVC "Controller LUN Number" becomes 0000000000000012
  - A9000/R capacity is displayed in Decimal whereas SVC uses Binary

- **Licensing usable capacity**
  - Contact Shelly Howrigon for an exception to license A9000/A9000R by usable capacity, not effective capacity

SVC DH8 Adapter Options

Slot	Possible Cards	Slot	Possible Cards
1	FC (8Gbps or 16Gbps)	4	Compression accelerator
2	FC (8Gbps or 16Gbps) or 10Gbps Ethernet/FCoE	5	FC (8Gbps or 16Gbps) or 10Gbps Ethernet/FCoE
3	12Gbps SAS Expansion or 16Gbps FC	6	Compression accelerator

Notes:

- At least one adapter card is required
- Up to three 8Gbps Fibre Channel adapters (12 ports)
- Up to four 16Gbps Fibre Channel adapters (8 ports)
- Up to one 10Gbps Ethernet/FCoE adapter (4 ports)
- Using any of slots 4-6 requires second CPU and 32GB cache upgrade

SVC SV1 Adapter options

	Min	Max
16Gb quad Fibre Channel	0	4
10Gb quad iSCSI / FCoE/ IP Rep	0	1
12Gb SAS	0	1
Compression Accelerator	0	2

Bank 1	Bank 2	Bank 3
Slot 1: N/A	Slot 3: 10Gb or 16Gb	Slot 6: 10Gb or 16Gb
Slot 2: 12Gb SAS	Slot 4: 10Gb or 16Gb	Slot 7: 10Gb or 16Gb
N/A	Slot 5: 12Gb SAS or Compression	Slot 8: Compression

Spectrum Virtualize DH8 and SV1 – 16 Ports

	1	2	3	4	5	6	7	8	9
1	**Slot/Port**	**Port #**	**SAN**	**4-port Nodes**	**8-port Nodes**	**12-port Nodes**	**16-port Nodes**	**16-port SV1 Nodes**	**16-port SV1 Nodes**
2								Using Compression	Non Compression
3	S1P1	1	A T0	Host/Storage/Inter-node	Host/Storage	Host/Storage	Host/Storage	Inter-node*	Inter-node*
4	S1P2	2	B T1	Host/Storage/Inter-node	Host/Storage	Host/Storage	Host/Storage	Host/Storage or Replication**	Host/Storage or Replication**
5	S1P3	3	A T2	Host/Storage/Replication (Inter-node if no replication planned)	Host/Storage	Host/Storage	Host/Storage	Host/Storage	Host/Storage
6	S1P4	4	B T3	Host/Storage/Replication (Inter-node if no replication planned)	Host/Storage	Host/Storage	Host/Storage	Host/Storage	Host/Storage
7	S2P1	5	A T4		Inter-node*	Inter-node*	Inter-node*	Host/Storage or Replication**	Host/Storage or Replication**
8	S2P2	6	B T5		Inter-node*	Inter-node*	Inter-node*	Host/Storage	Host/Storage
9	S2P3	7	A T6		Host/Storage or Replication**	Host/Storage or Replication**	Host/Storage or Replication**	Host/Storage	Host/Storage
10	S2P4	8	B T7		Host/Storage or Replication**	Host/Storage	Host/Storage	Inter-node*	Inter-node*
11	S3P1	9	A T4			Host/Storage	Host/Storage	Host/Storage	Host/Storage
12	S3P2	10	B T5			Host/Storage or Replication**	Host/Storage or Replication**	Host/Storage	Host/Storage
13	S3P3	11	A T6			Inter-node* or Host Storwage	Inter-node* or Host Storage	Inter-node* or Host Storage	Inter-node* or Host Storage
14	S3P4	12	B T7			Inter-node* or Host Storwage	Inter-node* or Host Storage	Inter-node* or Host Storage	Inter-node* or Host Storage
15	S5P1	13	A T0				Host/Storage	Host/Storage	Host/Storage
16	S5P2	14	B T1				Host/Storage	Host/Storage	Host/Storage
17	S5P3	15	A T2				Host/Storage	Host/Storage	Host/Storage
18	S5P4	16	B T3				Host/Storage	Host/Storage	Host/Storage
19									
20	* localfcportmask			11 or 1111 if no replication	110000	110000110000	110000110000	0000110010000001	0000110010000001'
21	** remotefcportmask			1100	11000000	1001000000	1001000000	0000000000010010	0000000000010010'
22									
23	**Note:** Column C assumes an odd/even SAN port configuration. Modifications will have to made if other SAN connection schemes are used.								
24									
25	**Note:** These recommendations represent optimal configurations based on assigning specific ports to specific uses and aligning the internal allocation of hardware CPU cores and software I/O threads to those ports.								
26	Therefore, varying from these recommendations may result in unexpected consequences. In addition, configuring as recommended above will ensure the ability to replace nodes non-disruptively in the future.								
27									
28	**Note:** It is not recommended to zone host or storage ports to ports designated for Inter-node use or Replication use in the 8/12/16 port configurations and in no case should inter-node and replication use the same ports.								
29	This is to minimize any B2B credit exhaustion situations, due to long distance latencies introduced by replication for example, from tieing up buffers needed by hosts, storage or inter-node communications.								
30									
31	**Note:** With 12 or greater ports per node we highly recommend 4 ports be dedicated for inter-node traffic especially when high write data rates are expected as all writes are mirrored over these ports within an I/O Group.								
32									
33	**Note:** Although supported, dedicating certain ports specifically for hosts and other ports for storage is not recommended as it negates the full duplex capability of the fibre channel HBA ports.								

How many paths?

- **A9000 model 415 / 425**
  - 6 paths, ports 1 and 3 from each grid controller

- **A9000R model 415**
  - 2 Grid elements
    - 8 paths, ports 1 and 3 from each grid controller
  - 3 Grid elements
    - 12 paths, ports 1 and 3 from each grid controller
  - 4 Grid elements
    - 8 paths, port 1 from grid controllers 1 through 4, port 3 from grid controllers 5 through 8
  - 5 Grid elements
    - 10 paths, port 1 from grid controllers 1 through 5, port 3 from grid controllers 6 through 10
  - 6 Grid elements
    - 12 paths, port 1 from grid controllers 1 through 6, port 3 from grid controllers 7 through 12

How many paths?

- **A9000R model 425**
  - 2 Grid elements
    - 8 paths, ports 1 and 3 from each grid controller
  - 3 Grid elements
    - 12 paths, ports 1 and 3 from each grid controller
  - 4 Grid elements
    - 16 paths, ports 1 and 3 from each grid controller

A9000/A9000R Volume Size for SVC MDisks

- **Data reduction should be done at the A9000/A9000R**

- **Considerations for determining what size volume to create on A9000/A9000R and present to SVC are:**
  - Determine data reduction rate
    - When in doubt, use 2.5
  - Multiply data reduction times usable capacity
  - Divide by
    - Number of paths x 3 for A9000
    - Number of paths x 2 for A9000R
  - Maximum of 511 LUNs from any one system can be mapped to a SVC cluster

  - Example 1 A9000 with 57TB FlashSystem 900
    - Data reduction rate is 6.8:1
    - 57 X 6.8 = 388 / (6 * 3) = 21.5 Create 18 volumes of 21.5 TB each

  - Example 2 A9000R model 425, 4 grid element, 180TB FlashSystem 900
    - Data reduction rate is 3.8:1
    - 180 X 4 X 3.8 = 2736 / (16 * 2) = 85.5 Create 32 volumes of 85.5 TB each

A9000/A9000R Configuration Considerations

- **A9000/R Snapshot, thin provisioning, replication, LUN expansion are not supported**

- **No need to reserve Snapshot space on A9000/R**

- **Configuring A9000/R host connectivity for the SVC cluster**
  - Use host cluster definition - preferred -
    - Create an A9000 cluster - use default host type
    - Add each SVC node as a host in the cluster
    - Apply SVC WWPNs to corresponding host object
    - Map all volumes to the cluster
    - Advantages
      - Allows for easier logical administration if multiple SVCs
      - Performance statistics per 'host' – **see which SVC I/O groups are busy**

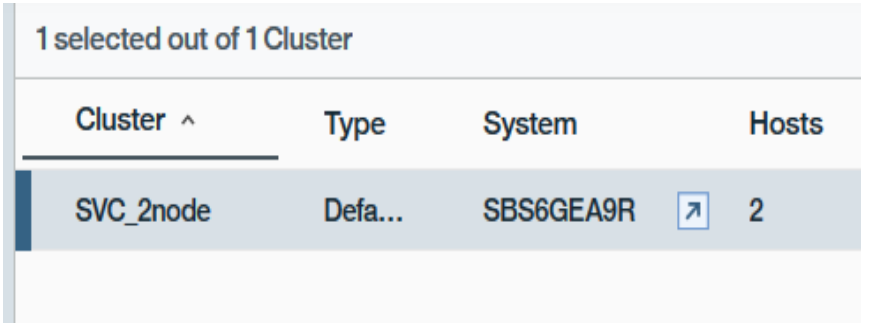

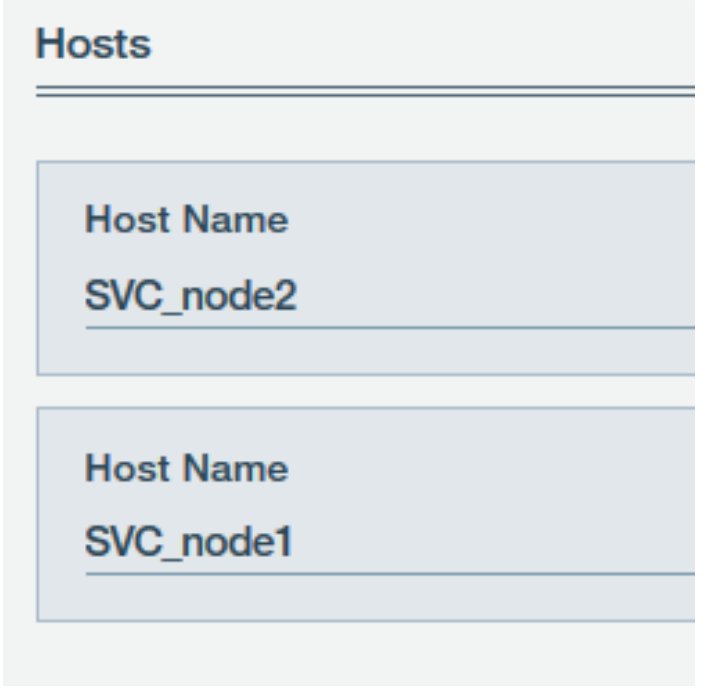

A9000/A9000R Configuration Considerations (continued)

- **Configuring A9000/R host connectivity for the SVC cluster**
  - Create one host definition on A9000/A9000R for SVC cluster - simpler
    - Create one host for the whole SVC cluster - use default host type
    - Apply SVC WWPNs to single host
    - Map all volumes to the host
    - Advantage
      - Easier to configure

A9000/A9000R patch to prevent OOP

- **Patch that can be applied by SSR at time of code load**
- **Reserves 300GB**
  - **Amount can be increased**
- **Creates notifications at 50%, 80%, 90% of physical capacity**
- **Takes machine offline if 100% of physical capacity is reached**
  - **Support can then release the 300GB, allowing for recovery/migration activity to occur**
- **Creates custom event if SVC host type is detected, and patch not applied, Service center opens a PMR**

Redbook: FlashSystem A9000 and A9000R, XIV and Spectrum Accelerate with IBM SAN Volume Controller Best Practices

Content:

- Chapter 1. IBM Spectrum Accelerate family and IBM System Storage SAN Volume Controller attachment
- Chapter 2. IBM XIV Gen3 with IBM System Storage SAN Volume Controller
- Chapter 3. IBM FlashSystem A9000 or IBM FlashSystem A9000R with IBM System Storage SAN Volume Controller
- Chapter 4. IBM Spectrum Accelerate with IBM System Storage SAN Volume Controller and IBM Storwize V7000
- Chapter 5. Data movement concepts: SAN Volume Controller and IBM Spectrum Accelerate family

New version

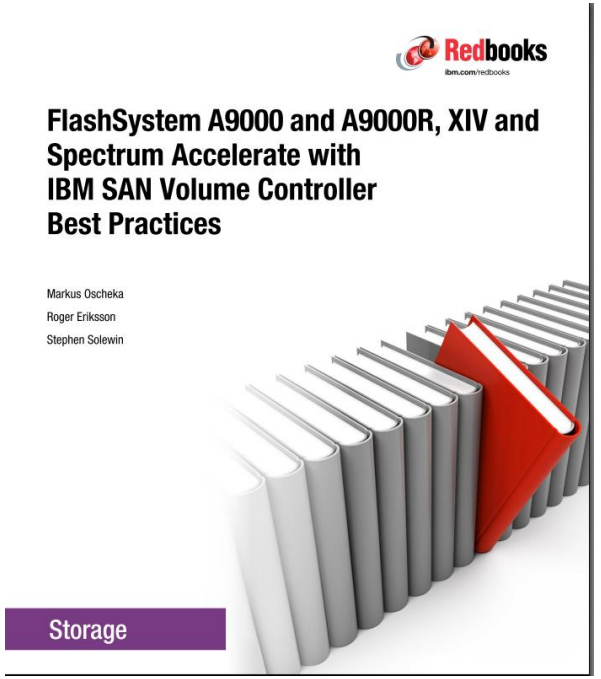

Links:
http://www.redbooks.ibm.com/abstracts/redp5408.html?Open

Download on the iBookstore

GET IT ON Google play

Thank You

Questions

Where do you find more information?

- A9000R Knowledge Center: https://www.ibm.com/support/knowledgecenter/en/STJKN5/landing/IBM_FlashSystem_A9000R_welcome_page.html
- Fix Central for IBM FlashSystem™ A9000R
- A9000 Knowledge Center: https://www.ibm.com/support/knowledgecenter/en/STJKMM/landing/IBM_FlashSystem_A9000_welcome_page.html
- Fix Central for IBM FlashSystem™ A9000
- Spectrum Control Base Edition Knowledge Center https://www.ibm.com/support/knowledgecenter/en/STWMS9/landing/IBM_Spectrum_Control_Base_Edition_welcome_page.html
- A9000 Technical Support Information newsletter: Link
  - Current issue: http://www-01.ibm.com/support/docview.wss?uid=ssg1S7005397